

REAL-TIME SEARCH ENGINE

By:

Shawn Fanning

John Fanning

Edward Kessler

Background

Search engines as they currently exist on the Internet are used by people all over the world to find and download data objects of interest that reside on servers. Typically, these search engines periodically examine many servers on the Internet to see what data objects each server contains. Thereafter, the search engine constructs an index of each server's contents, and links the contents to that server's location.

The construction of the index is a time consuming task, and because of the relative cost involved to the servers and the search engine, it cannot be done very often. The timeliness of the information created by the search engine is sacrificed in order to reduce the burden on the index builder of the search engines and the servers that contain the data being searched.

This means that the search engine index is quickly out of date. For some types of data objects, this matters very little, since the data objects are created and modified relatively slowly. However, for data objects that are created and removed relatively often, the search engine indices are impractical, and for data objects that are added and removed daily, the standard search engines are practically useless. In addition, the current paradigm assumes a relatively static server environment, but in an environment where servers come up and go down relatively frequently and data objects are added and deleted hourly or more frequently, the standard search engine methodology is not useful at all.

Thus, it can be seen that there is a need for an Internet search engine that maintains an up-to-date index of data content residing on servers that are currently

connected to the Internet.

There is a further need for a real-time search engine that significantly reduces the cost of constructing a search engine index using methods employed by the prior art.

Summary

5 The present invention satisfies these needs by providing a method for creating a real-time search engine over the Internet that provides a search response containing data object descriptions and server descriptions 34 of data objects that are currently available for transfer from a provider server directly to a recipient client in response to a recipient client search request. This method comprises the provider server connecting to a real-
10 time search engine through the Internet, the provider server providing the real-time search engine with data object descriptions of data objects residing on the provider server, and the real-time search engine indexing data object descriptions associated with the data object of the provider server, wherein the data object descriptions provided by the provider server are purged from the real-time search engine when the provider server
15 is disconnected from the real-time search engine.

The method further comprises the provider server automatically, in real-time, providing the real-time search engine with data object descriptions of data objects that are added to the provider server.

The method preferably further comprises the provider server automatically, in
20 real-time, notifying the real-time search engine of data objects that are removed from the provider server, wherein the real-time search engine then purges the data object descriptions.

The data object descriptions comprise any of the following: a title of the data object, the size of the data object, the type of data object, any text associated with the data
25 object, the creator of the data object, the quality rating of the data object, and the provider server on which the data object resides. The server description 34 comprises any of the following: the server Internet Protocol address, the number of simultaneous connections allowed by the server, the server's reliability, and the server's name.

Preferably, a client search command is used, wherein a recipient client searches

the data object descriptions to find the best data object and selects the most optimal provider server that the data object resides on.

Furthermore, the recipient client search request further comprises a provider server limitation criteria, wherein the search engine prunes the search response of all provider servers that do not meet the server limitation criteria.

In a preferred embodiment, the provider server limitation criteria comprises a bandwidth limitation, wherein the search engine prunes the search response of provider servers that have a bandwidth capability that is below the bandwidth limitation.

Optionally, the real-time search engine purges from the search response provider servers that cannot accept additional recipient client download requests.

Also in a preferred embodiment, an automated search response sort by the client. The automated search response is sorted by the responsiveness value, wherein the responsiveness value is determined by measuring the amount of time an echo reply message takes to be returned by the provider server to the recipient client. Preferably, the provider server is pruned from the search response if the provider server did not respond to the recipient client's echo request within a specified period of time.

The data object is of the type selected from the group comprising: an audio data object, a text data object, a image data object, a video data object, and a software executable data object.

In a preferred embodiment, the real-time search engine further comprises the recipient selecting one of the provider servers in the search response, and then the recipient client downloading the data object from the selected provider server. Additionally, the recipient client simultaneously operates as a provider server to other recipient clients, making data objects that have been downloaded by the recipient client available to other recipient clients on the Internet.

Brief Description of the Drawings

Fig. 1 is an overview diagram of a preferred embodiment of the system of the present invention;

Fig. 2 is an overview diagram of a preferred embodiment of the real-time search engine with its search engine, index builder and gateway components;

Fig. 3 is a flow chart of an embodiment of the process when a new data object is downloaded from the Internet, or otherwise added to the provider server, or during the initial scan of the data object collection during the log-in process; and

Fig. 4 is a flow chart of an embodiment of a search request constructed by a recipient client.

Detailed Description

The invention provides a method for creating a real-time search engine over the Internet that provides a search response containing data object descriptions and server descriptions of data objects that are currently available for transfer from a provider server directly to a recipient client in response to a recipient client search request. It is the essence of this invention that data object descriptions provided by the provider server are purged from the real-time search engine when the provider server is disconnected from the real-time search engine. This provides the recipient client with search results that include only those server sources that are currently available to provide and transfer any data to the recipient client.

Turning now to the figures, the overall configuration of the invention and its components are shown in Fig. 1. Essentially the components of a real-time search includes at least one recipient client 16 which connects to the real-time search engine 10 to locate a desired data object 24. At least one provider server 12 connects to the real-time search engine and provides one or more data object descriptions 22 to the real-time search engine. If the provider server 12 is disconnected from the real-time search engine, the data object descriptions 22 are purged from the real-time search engine, indicating that those data objects 24 are no longer available for download from the provider server to the recipient client 16.

Preferably, the provider server automatically, in real-time, provides the real-time search engine with data object descriptions 22 of data objects 24 that are added to the provider server 12.

Also, preferably, the provider server also automatically, in real-time, notifies the real-time search engine 10 of data objects 24 that are removed from the provider server 12, wherein the real-time search engine then purges the data object descriptions 22.

Recipient Client

5 Recipient clients connect to the real-time search engine 10 to find the best provider server that contains the particular data object that the recipient client 16 wishes to download from a provider server. The recipient client preferably uses a recipient browser 18 for communicating with the real-time search engine 10 and for making search requests from the real-time search engine. Examples of these browsers include the
10 Netscape Communicator or the Microsoft Explorer or other custom interfaces.

Connections are typically software protocols that provide a method for transmitting information between entities that are connected; an example of such a protocol is TCP, which is the preferred connection protocol for the invention. However, other protocols that fulfill the same basic functionality as TCP (such as a UDP protocol
15 with retransmits, and a disconnection timeout) will also suffice. These protocols are well known in the art.

In another embodiment, where the data object 24 is a text file, the search request contains any one of the following: a partial filename, keywords, author, the size of the file, the category, and the description of the text.

20 In one embodiment, where the data object 24 is an audio data file, the search request contains any one of the following: a partial filename, a bitrate (bps), a sample rate (Hz) of the data, the size of the file, the duration, the name of the author or artist, the song title, the genre, and the title of the album.

25 In another embodiment, where the data object 24 is an image or video file, the search request contains any one of the following: a partial filename, the amount and type of compression, the size of the file, the category, and the description of the image or video.

When a search response is returned to the recipient, the recipient browser 18 displays the results of the search request for the recipient to examine.

In a preferred embodiment, the recipient client 16 determines a response time of each of the provider servers returned in the search response. The response time is measured by the recipient client 16 sending an ICMP echo packet to each provider server, and measuring the amount of time it takes to receive a reply from the provider server.

- 5 The recipient client browser 18 uses the response time to sort the data object descriptions 22, which then displays the data object descriptions of the provider servers in order of their response time.

- In an embodiment, the recipient may choose a search parameter for the real-time search engine 10 to provide a search response 38 that includes only data object descriptions on provider servers that have a minimum data transfer bandwidth capability.

In another embodiment, the recipient directs the search engine to return a search response 38 containing only data object descriptions for provider servers that are not currently too busy to accept additional download requests.

- In one embodiment, the provider server is not located behind a firewall. The recipient client 16 downloads a data object 24 from the provider server by connecting directly to the provider server, requesting a data object, and then storing the data object in the recipient's data object collection.

- In one embodiment, an optimal provider server is automatically selected from among at least two provider servers that are able to supply a desired data object using a scoring mechanism. The scoring mechanism comprises the roundtrip response time from the recipient client to the provider server, the Internet connection line speed (data transfer speed) of the provider server, the size of the file, and the reliability of the provider server. The best score is usually from a provider server that has a high line speed and high provider server reliability. The provider server with the best score is preferably selected by the recipient client for download.

In another embodiment, in order to determine the best score, the recipient client or the provider server uploads to the real time search engine the actual transfer rate for each data object transfer which is used to calculate of the effective line speed of the provider server.

Provider Server

Each provider server contains a data object collection of data objects **24** that may be downloaded from the provider server. When the provider server is prepared to provide
5 data objects to any requesting recipient client **16**, the provider server connects to the real-time search engine, and uploads descriptions of each data object in the data object collection. The real-time search engine is updated immediately. The data object descriptions **22** comprise any of the following: a title of the data object, the size of the data object, the type of data object, any text associated with the data object, the creator of
10 the data object, the quality rating of the data object, and the provider server on which the data object resides.

In the preferred embodiment, the connection between provider server and real-time search engine **10** is accomplished using the TCP protocol. Occasional messages are sent between provider server **12** and the real-time search engine to assert that the
15 connection between the two is valid. If no message is received from the provider server for several minutes, the connection is closed and the connection to the provider server is broken.

In one embodiment, the provider server authenticates itself to the real-time search engine using a login process, immediately after connecting to the real-time search engine,
20 by transmitting a login name and a password.

In another embodiment, a determination is made if the provider server **12** is protected by a firewall, and this determination is transmitted to the real-time search engine **10** during the initial login.

In yet another embodiment, when the provider server scans the data objects in the
25 data object collection, each data object's type is ascertained by examining the extension on the filename (.mp3, .jpg, .mpg, .doc are a few examples). Files without extensions are ignored. Each file is validated as to the proper formatting of the data contained within. Data objects that fail validation do not have their descriptions uploaded to the real-time search engine.

When data objects are added to the provider server, the provider server transmits the new data object's description to the real-time search engine. Likewise, when a data object 24 is deleted, the provider server 12 notifies the real-time search engine of the deletion.

5 In one embodiment, during the login process the provider server only transmits the changes that were made in its data object collection since the last connection to the real-time search engine. Both the real-time search engine 10 and the provider server store a copy of the data object descriptions that have been uploaded to the real-time search engine, and all of the successfully acknowledged changes to those descriptions. In this way, the initial information transmitted from the provider server to the real-time search engine is minimized for large data object collections.

In the preferred embodiment, the data object collection is at least one directory on the provider server. The data object collection alternatively contains other directories that themselves contain other data objects or more directories.

15 In another embodiment, the data object collection is stored on a computer remote from the provider server 12, but is accessible by the provider server. A data object collection is optionally data objects in a database, files in a directory, data objects in memory, on CD-ROM, flash memory, etc.

20 In one embodiment, the provider server also contains a server description, which comprises its own data transfer line bandwidth to the Internet, and it uploads this server description during the initial connection to the real-time search engine.

In a preferred embodiment, both the provider server and recipient client 16 are located within the same executable image. Thus, whenever a recipient runs a recipient browser, he also simultaneously runs a provider server.

25 In one embodiment, data objects downloaded by the recipient client from other provider servers are immediately added to the data object collection, making these data objects instantly available to other recipient clients on the Internet. In this embodiment, the rapid spread of data objects throughout the network of provider servers and recipient clients is greatly facilitated.

In a preferred embodiment, a data object fingerprint is constructed by performing a checksum of the data object. Each data object is uniquely identifiable by the fingerprint of the data object's data.

In a preferred embodiment, if the provider server 12 is not behind a firewall, recipient clients connect directly to the provider server, and request that a chosen data object be transferred from the provider server and downloaded to the recipient client 16. If the provider server is behind a firewall, then the recipient client 16 asks the real-time search engine 10 to pass the download request to the provider server. When the provider server receives this download request, it then connects to the recipient client 16 and then the download occurs. If both the provider server 12 and the recipient client are protected by firewalls, a proxy server is used to facilitate the transfer. The recipient client informs the real-time search engine of the download request, the real-time search engine transmits the request to the provider server, the recipient client and the provider server both connect to the proxy server, which then allows data to flow and hence the download to occur between the recipient client and the provider server through the proxy server.

Real-Time Search Engine

As shown in Fig. 2, in a preferred embodiment, the real-time search engine has the following components: at least one search engine Gateway, at least one search engine, at least one search engine Index Builder, and a search engine Index.

Search Engine Gateway

In the preferred embodiment, each provider server that connects to the real-time search engine connects to the search engine gateway 44 component. When a provider server uploads information about the data objects it contains, the search engine gateway 44 takes this information and passes it to the search engine index builder 42, which uses it to update the search engine index. When an provider server disconnects, or is disconnected by a network error, or otherwise fails to communicate with the real-time search engine, the search engine gateway detects this, and informs the search engine index builder 42, which in turn removes the data object descriptions uploaded by that provider server 12 from the search engine index.

Alternatively, when a particular provider server is disconnected, the search engine index builder 42 does not actually remove the data object, but instead marks the data object descriptions as "Not Available." When that provider server re-connects, instead of transmitting the entire list of data object descriptions, it transmits only changes to its data object collection that may have occurred during the disconnected period. During searches, the search engine 40 ignores all data object descriptions that are marked as "Not Available."

In one embodiment, each recipient client 16 also connects to a search engine gateway. In this embodiment, each search engine gateway 44 connects in turn to a search engine 40. All search requests from recipient clients are transmitted to the search engine gateway, and the search engine gateway then transmits the search requests to the connected search engine. The search engine executes the search request, and transmits the search response 38 back to the search engine gateway, which in turn transmits the search response back to the originating Recipient client.

In another embodiment, the search engine gateway tracks data object downloads initiated by recipient clients. The recipient client transmits a request to download a particular data object from a provider server. If the download is successful, the recipient client 16 informs the search engine gateway 44 that the download was completed. Using this information, the search engine gateway tracks the reliability of the provider server, as well as the current number of recipient clients downloading data objects from a particular provider server.

Search Engine

The search engine receives search requests 36 from recipient clients, executes the search requests, and constructs search responses. The search responses are transmitted back to the recipient clients.

In another embodiment, the search engine also receives search requests from search engine gateways, that are simply relaying the search requests from recipient clients.

In the preferred embodiment, a Search request contains: a partial data object

name, an optional minimum data object quality rating, an optional minimum provider server connection bandwidth, and an optional maximum number of data object descriptions to be retrieved.

Each search response contains a list of data object descriptions as well as a list of server descriptions. In the preferred embodiment, a subset of the fields in the data object descriptions and server descriptions are returned in the search response, including: a provider server name and Internet Protocol (IP) Address, a provider server bandwidth description (56k modem, DSL, T1, etc), a data object name (in the audio embodiment, the song title and artist name), a data object fingerprint, a data object size (in bytes), and a data object quality rating.

To execute the search, the search engine uses the fields in the search request to scan the records stored in the search engine index. If a particular data object description is marked as "Not Available" it is ignored by the search engine.

In one embodiment, the search engine gateway and the search engine exist in the same process. In another embodiment, the search engine gateway and the search engine exist on different processes, but run on the same machine. Many configurations of machines, search engine gateways, and search engines are possible.

Search Engine Index

In a preferred embodiment, the search engine index has two internal tables. These tables include a data object description table and a provider server description table. These tables are managed by the search engine index builder.

The provider server description table contains a collection of provider server descriptions. Some of these fields are uploaded by the provider servers during the initial connection to the search engine gateway. Others are calculated as events occur. In the preferred embodiment, entries in this table contain the following fields:

- server name & IP Address
- password
- connection bandwidth (T1, 56k modem, DSL, etc)
- must push data objects to recipient client?

- list of data object descriptions for this server
- remaining available connections allowed by Provider server
- site reliability (% of successful transfers)

The data object description table contains a collection of data object descriptions

22. These are uploaded by the provider server 12. As data objects are added, new data object descriptions are uploaded. As data objects are removed, existing data object descriptions are removed or optionally marked for removal. In a preferred embodiment, entries in this table contain the following fields:

- filename
- metadata (in the audio embodiment, song name, artist name, song description)
- the data object fingerprint
- size (in bytes)
- quality rating (in the audio embodiment, the encoding bit rate and sampling frequency)
- a link to the Provider server Description record

Execution Processes

Five different execution processes are serviced by the system: data object added, data object removed, search request, provider server connect, provider server disconnect.

Data Object Added

In an embodiment as shown in Fig. 3, when a new data object is created on a provider server, when a new data object is downloaded from the Internet, or during the initial scan of the data object collection during the log-in process, the following processes occur:

- a) the data object fingerprint is calculated,
- b) the data object is given a name, a source name, size, and a quality rating, forming an data object description,
- c) the data object description is uploaded to the search engine Gateway, and passed to the search engine index builder,
- d) The search engine index builder creates a new data object description entry in

the search engine index, and

e) the search engine Index Builder updates the Provider server Description entry for this server to reflect the new data object description Entry.

Data Object Removed

In an embodiment, when an existing data object is removed, the following occurs:

- a) the data object fingerprint is retrieved,
- b) the data object fingerprint is transmitted to the search engine gateway, and passed to the search engine index builder,
- c) the search engine index builder removes the data object description entry for that provider server, and
- d) the search engine index builder updates the provider server description entry for that provider server to reflect the removal of the data object description.

Search Request

In an embodiment as shown in Fig. 4, when a search request is constructed by a recipient client 16, the following occurs:

- a) the search request is uploaded to the search engine,
- b) the search engine searches the name column of the data object description table for all matches on the data object name,
- c) the search engine prunes the resulting data object description list using the provider server bandwidth limitation and the minimum quality rating limitation,
- d) if at any time the number of data object descriptions returned exceeds the maximum number of data object descriptions limitation, the search terminates and no further data object descriptions are retrieved, and
- e) the resulting list of data object descriptions and related server descriptions are sent to the recipient client.

Provider Server Connect

In an embodiment, when a provider server first connects to the real-time search engine, the following occurs:

- a) a provider server description record is created for the provider server,

- b) data object descriptions for all data objects in the provider server's data object collection are uploaded to the search engine gateway, and passed to the search engine index builder, and
- c) the search engine index builder treats each uploaded data object description as a data object added process.

Provider Server Disconnect

In an embodiment, when an provider server disconnects from the real-time search engine, the following occurs:

- a) the search engine index builder removes all data object descriptions referring to this Provider server as in the data object removed process, and
- b) the search engine index builder deletes or optionally marks for deletion the provider server description record.

Alternate Embodiments

In one embodiment, the data objects are audio files, and the data object descriptions comprise the filename, the bit rate, sampling frequency, and size obtained from the audio file itself. In this embodiment, preferably the recipient client 16 also incorporates an audio player, for playing the audio file. Also, the provider server contains a mechanism for constructing an audio file from a CD or other audio media source, that deposit newly-created audio files into the data object collection.

In another embodiment, the data objects are image and video files, and the data object descriptions include the filename, the compression detail and other information obtained from the jpg file itself, as well as a short description of the image. In this embodiment, preferably the recipient client 16 application also incorporates a means for displaying the image or video file, and the provider server incorporates a means for generating an image or video file from a photo or other visual image source.

In yet another embodiment, the data objects may be text, audio, image, and video data objects. Example formats include HTML text, MP3 audio, JPEG still image, and MPEG video data. Each different type of data object is also then distinguishable by type, as well as by name, and the other attributes mentioned previously.

As new image sources, and image compression and storage mechanisms become available, data object generation methods for these protocols and storage formats can be added to the recipient client and provider server without deviating from the spirit of this invention.